

Serial No. 09/940,728
Page 13 of 25

REMARKS

This response is intended as a full and complete response to the non-final Office Action mailed May 27, 2003. In the Office Action, the Examiner notes that claims 1-39 are pending, of which claims 1-39 are rejected. By this amendment, Applicant has amended claims 1-10, 16-17, 23, 25, and 28, canceled claim 24, and claims 11-15, 18-22, 26-27, and 29-39 continue unamended.

In view of both the amendments presented above and the following discussion, Applicant submit that none of the claims now pending in the application are non-enabling or anticipated under the respective provisions of 35 U.S.C. §112 or §102. Thus, Applicant believes that all of these claims are now in allowable form.

It is to be understood that the Applicant, by amending the claims, does not acquiesce to the Examiner's characterizations of the art of record or to Applicant's subject matter recited in the pending claims. Further, Applicant is not acquiescing to the Examiner's statements as to the applicability of the art of record to the pending claims by filing the instant responsive amendments.

IN THE SPECIFICATION:

The Applicant has amended the specification to provide minor grammatical corrections, change reference designations to conform to the reference designations in the drawings, and correct the title on the cover sheet of the application to conform to the title on the first page of the application. Further, the Applicant has amended the specification to update serial numbers associated with patent applications that have been incorporated by reference therein to their corresponding patent numbers. Such grammatical corrections, reference designation changes, and changes to the incorporated reference designations do not add any new subject matter to the application.

CLAIM OBJECTIONS

The Examiner has objected to claim 23 stating: "Claim 23 recites the limitation 'parity group containing the failed disk drive'. Applicant respectfully traverses the Examiner's objection.

279205-1

Serial No. 09/940,728
Page 14 of 25

In response to the Examiner's objection, Applicant has amended claim 23 to recite in part: "parity group containing an extent associated with the failed disk drive." As such, the Applicant submits that the feature is clearly defined and no new subject matter has been added. Therefore, the Applicant respectfully requests that the objection be withdrawn.

REJECTIONS

35 U.S.C. §112

Claims 8, 16 and 17

The Examiner has rejected claims 8, 16 and 17 under 35 U.S.C. 112, ¶1, as failing to comply with the written description requirement. Applicant respectfully traverses the Examiner's rejection.

With respect to claim 8, the Applicant has amended claim 8 to recite features with respect to the second parity segment. In particular, claim 8, as amended, recites:

"The method of claim 1 wherein said generating said second parity segment comprises computing second parity information from each first parity segment in said parity group."

Since the scope and features of claim 8 has been changed, the rejection is now considered moot. Therefore, the Applicant respectfully requests that the rejection be withdrawn.

With respect to claim 16, the Applicant has amended the feature "at the end of one of said m-1 data extents" to "at the end of one of said m data extents." Support for such change may be found in Applicant's FIG. 4, where 1/m extents includes the second parity segment.

As such, the Applicant submits that claim 16, as amended, fully satisfies the requirements under 35 U.S.C. 112, ¶1, and is fully patentable thereunder. Furthermore, claim 17 depends from claim 16 and recites additional features thereof. As such and for at least the same reasons discussed above, the Applicant submits that this dependent claim also fully satisfies the requirements under 35 U.S.C. 112, ¶1, and

279205-1

Serial No. 09/940,728
Page 15 of 25

is fully patentable thereunder. Therefore, the Applicant respectfully requests that the rejections be withdrawn.

Claims 2, 4-9, and 28-39

The Examiner has rejected claims 2, 4-9, and 28-39 under 35 U.S.C. §112, ¶2, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. Applicant respectfully traverses the rejection.

Claims 2 and 4-6

With respect to claims 2 and 6, and in particular the limitation "said parity segments," the Examiner states that it is not clear whether "said parity segments" refers to the first mentioned (in the parent claim) "parity segments" only or includes the "second parity segment". The Applicant respectfully traverses the rejection.

The Applicant has amended independent claim 1 to properly distinguish the types of parity segments. Specifically, claim 1 recites in part:

"generating first parity segments from said data segments associated with each parity group;
storing said first parity segments;
generating a second parity segment from said first parity segments."
(emphasis added).

Furthermore, the Applicant has amended claims 2-4 and 6, which depend, either directly or indirectly, from independent claim 1 and recite additional features thereof. In particular, the Applicant has amended these dependent claims to further clarify and identify the first and second parity segments associated with the Applicant's invention. Accordingly, the Applicant submits that these dependent claims, as amended, particularly point out and distinctly claim the subject matter which applicant regards as the invention. As such, the Applicant submits that these claims fully comply with the requirements under 35 U.S.C. §112, ¶2, and are patentable thereunder. Therefore, the Applicant respectfully requests that the rejections be withdrawn.

279205-1

Serial No. 09/940,728
Page 16 of 25

Claim 5

With respect to claim 5, referring to the language "said parity segments step further comprises the step ...," the Examiner states that "it seems that this claim is either combining the two previously claimed steps [of the parent claim] or adding another step of storing the second parity segment." The Applicant respectfully traverses the rejection.

The Applicant has amended claim 5 to clarify the features the Applicant considers as being inventive. In particular, claim 5, as amended, recites:

"The method of claim 4 wherein said storing said second parity segment comprises appending said second parity segment after said data segments in the nth extent." (emphasis added).

The Applicant submits that the features of claim 5, as amended, now properly identify and particularly point out and distinctly claim the subject matter which applicant regards as the invention. As such, the Applicant submits that claim 5 fully comply with the requirements under 35 U.S.C. §112, ¶2, and are patentable thereunder. Therefore, the Applicant respectfully requests that the rejection be withdrawn.

Claims 7 and 9

With respect to claims 7 and 9, the Examiner states that "[t]here are two generating steps in the claim" and "[i]t is not clear which one of the two 'generating steps' this limitation [said generating step] refers to." The Applicant respectfully traverses the rejection.

Claims 7 and 9 have been amended to further clarify the features the Applicant considers inventive. In particular, claims 7 and 9, as amended, recite:

7. The method of claim 6 wherein said generating said first parity segments comprises computing, for each first parity segment, parity information associated with a data segment formed contiguously across each extent of said parity group.

9. The method of claim 8 wherein said second parity segment comprises parity information associated with all of the data segments in said parity group.

279205-1

Serial No. 09/940,728
Page 17 of 25

Since the scope and features of claims 7 and 9 have been changed, the rejections are now considered moot. Therefore, the Applicant respectfully requests that the rejections be withdrawn.

Claims 28-39

The Examiner states that there is insufficient antecedent basis for the plural version of the feature "user access to the at least one file" in these claims and suggests that it would make more sense if claim 28 were dependent on claim 27, rather than claim 7.

The Applicant has amended claim 28 to correct the previously mistyped dependency from claim 7 to now depend from claim 27. Accordingly, amended claim 28 now depends indirectly (via claim 27) from independent claim 23, which provides sufficient antecedent basis by reciting the feature "the at least one file."

As such, the Applicant submits that claim 28 has proper antecedent basis for all the limitations therein and fully complies with the requirements under 35 U.S.C. § 112, and is patentable thereunder. Furthermore, claims 29-39 depend from claim 28 and recite additional features thereof. As such and for at least the same reasons discussed above, the Applicant submits that these dependent claims have proper antecedent basis for all the limitations therein and fully comply with the requirements under 35 U.S.C. § 112, and are patentable thereunder. Therefore, the Applicant respectfully requests that the rejections be withdrawn.

35 U.S.C. §102

Claims 1-22

The Examiner has rejected claims 1-22 under 35 U.S.C. §102(e) as being anticipated by Han et al. (U.S. Patent 6,158,017, hereinafter "Han"). Applicant respectfully traverses the rejection.

Independent claim 1 (claims 10, 12 and 18 recite similar limitations as recited in independent claim 1) recites features of the Applicant's invention that the Applicant considers to be inventive. In particular, independent claim 1 recites:

" A method for storing data on an array of disk drives, comprising:

279205-1

Serial No. 09/940,728
Page 18 of 25

dividing at least a portion of storage capacity of each disk drive of said array into commonly sized extents for storing data and parity information;
dividing each extent into a plurality of data segments;
 defining a parity group comprising at least three contiguous extents formed across said array of disk drives;
 storing said data in data segments associated with each parity group;
generating first parity segments from said data segments associated with each parity group;
 storing said first parity segments;
generating a second parity segment from said first parity segments; and
 storing said second parity segment." (emphasis added).

"Anticipation requires the presence in a single prior art reference disclosure of each and every element of the claimed invention, arranged as in the claim" (Lindemann Maschinenfabrik GmbH v. American Hoist & Derrick Co., 730 F.2d 1452, 221 U.S.P.Q. 481, 485 (Fed. Cir. 1984)(citing Connell v. Sears, Roebuck & Co., 722 F.2d 1542, 220 U.S.P.Q. 193 (Fed. Cir. 1983)) (emphasis added). Han fails to disclose each and every element of the claimed invention, as arranged in the claim. Specifically, Han fails to teach the features "dividing each extent into a plurality of data segments," "generating first parity segments from said data segments associated with each parity group," and "generating a second parity segment from said first parity segments."

In particular, Han discloses if a disk array includes N (prime number) disks, and each of the disks is logically divided into N-1 blocks, then it is defined as a matrix $a(ij)$ of $(N-1)*N$, where $0 \leq i \leq N-2$ and $0 \leq j \leq N-1$. Matrix $a(ij)$ (as shown in TABLE 1 of Han) represents the data value of the logical block i of disk j, which contains a binary Number therein. (see Han, col. 5, lines 20-44). Nowhere in the Han reference is there any teaching that these logical blocks are further divided into segments, as recited in the Applicant's invention.

Specifically, the Applicant's invention provides:

FIG. 4 depicts a RAID 3+5 data striping format across a plurality of extents in the disks 320 in a single parity group 404. The parity group 404 is illustratively depicted as having four data extents 406 (extents 0-3) containing data (e.g., video information). Each data extent 406 is divided into segments 402₁ to 402₃ (collectively segments 402). The number of segments 402 in each extent containing data is equal to P-1, where P equals the number of disks 320 in a parity group 404. For example, the true (non-parity) data of extent 0 110₀ is divided into three segments (a, b and c). The true data of extent 1 110₁ is divided

279205-1

Serial No. 09/940,728
Page 19 of 25

into segments d, e, and f. The true data of extent 2 110₂ is divided into segments g, h, and i, and the true data of extent 3 110₃ is divided into segments j, k, and l.

The parity column 408 is a representation of the parity information for the segments as the data is stored across the parity group 404. For instance, the parity data "X" for the first row of extent segments 402₁ includes the data in segments a, d, g, and j of extents 0-3, respectively. Similarly, the parity data "Y" for the second row of extent segments 402₂ includes the data in segments b, e, h, and k of extents 0-3, respectively. Furthermore, the parity data "Z" for the third row of extent segments 402₃ includes the data in segments c, f, i, and l of extents 0-3, respectively. It should be noted that the parity information in column 408 is not a physical disk drive 320 that stores the parity information. Rather, column 408 is simply a representation of the parity information for the four data extents 406 of the parity group 404.

At the bottom of extent 0 110₀ is appended the parity segment X 410₀, which contains the parity data of the first segments 402₁ of each extent ($X = a \wedge d \wedge g \wedge j$) in the parity group 404, where " \wedge " indicates a Boolean logic "exclusive OR" (XOR) operative. Similarly, at the bottom of extent 1 110₁ is appended the parity segment Y 410₁, which contains the parity data of the second segments 402₂ of each extent ($Y = b \wedge e \wedge h \wedge k$) in the parity group 404. Likewise, at the bottom of extent 2 110₂ is appended the parity segment Z 410₂, which contains the parity data of the third segments 402₃ of each extent ($Z = c \wedge f \wedge i \wedge l$) in the parity group 404. Furthermore, at the bottom of the last extent 110₃ (i.e., extent 3) is appended the parity segment "Q" 410₃, which contains the parity information of the 3 parity segments X, Y, and Z, where $Q = X \wedge Y \wedge Z$. Thus, in the inventive RAID 3+5 format, the "pure parity" disks required for RAID 3 (see FIG. 1) are eliminated from the disk drive array 319. (see Applicant's specification, pages 9 and 10, paragraphs 0033, 0034, and 0036 and FIG. 4).

Accordingly, Han fails to disclose Applicant's claimed "dividing each extent into a plurality of data segments," "generating first parity segments from said data segments associated with each parity group," and "generating a second parity segment from said first parity segments." Therefore, Han fails to teach each and every element of the claimed invention, as arranged in the claim.

As such, Applicant submits that independent claim 1 is not anticipated and fully satisfies the requirements of 35 U.S.C. §102 and is patentable thereunder. Independent claims 10, 12, and 18 recite similar limitations as recited in independent claim 1. As such, and for at least the same reasons as discussed above, the Applicant submits that independent claims 10, 12 and 18 also are not anticipated and fully satisfy the requirements of 35 U.S.C. §102 and are patentable thereunder. Furthermore, claims 2-9, 11, 13-17 and 19-21 depend, either directly or indirectly, from independent

279205-1

Serial No. 09/940,728
Page 20 of 25

claims 1, 10, 12 and 18 and recite additional features thereof. As such and at least for the same reasons as discussed above, Applicant submits that these dependent claims are also not anticipated and fully satisfy the requirements under 35 U.S.C. §102 and are patentable thereunder. Therefore, Applicant respectfully requests that the rejections be withdrawn.

Claim 23

The Examiner has rejected independent claim 23 under 35 U.S.C. §102(e) as being anticipated by Allegrezza (U.S. Patent Pub. No. 2002/0157113, hereinafter "Allegrezza").

The Applicant has amended independent claim 23 to include the features of dependent claim 24. Applicant's independent claim 23 recites:

"A method for recovering data from a failed disk drive in a disk array comprising a plurality of disk drives, said plurality of disk drives having a plurality of extents apportioned equally into a plurality of parity groups, said method comprising:

striping data from at least one file sequentially across the plurality of extents of the plurality of parity groups, said striping data comprising dividing each extent in each parity group into a plurality of data segments, storing data information in the plurality of data segments, forming a plurality of parity segments from the data segments in the parity group, and appending one parity segment to the end of each extent, wherein parity segments in the parity group respectively store parity information corresponding to the parity group in which the parity segments reside;

identifying a failed parity group containing an extent associated with the failed disk drive;

outputting, in realtime to users in a normal disk access mode, said data from parity groups without the failed disk drive;

reconstructing, in a parity correction mode of operation, said data from the at least one failed parity group; and

outputting, in realtime to the users in the failed parity group, said reconstructed data." (emphasis added).

Allegrezza discloses reconstructing data for a failed drive, and after the data is reconstructed, reading and transmitting may continue as normal (see Alegrezza, page 3, paragraphs 0043-0046). Nowhere in Allegrezza is there any teaching, or even suggestion of "said striping data comprising dividing each extent in each parity group

279205-1

Serial No. 09/940,728
Page 21 of 25

into a plurality of data segments, storing data information in the plurality of data segments, forming a plurality of parity segments from the data segments in the parity group, and appending one parity segment to the end of each extent, wherein parity segments in the parity group respectively store parity information corresponding to the parity group in which the parity segments reside." In fact, the Allegrezza reference is completely silent with respect to dividing each extent into data and parity segments. Therefore, Allegrezza fails to teach each and every element of the claimed invention, as arranged in the claim.

As such, Applicant submits that independent claim 23 is not anticipated and fully satisfies the requirements of 35 U.S.C. §102 and is patentable thereunder. Therefore, Applicant respectfully requests that the rejection be withdrawn.

35 U.S.C. §103

Claims 23-27

The Examiner has rejected claims 23-27 as being obvious and unpatentable under the provisions of 35 U.S.C. §103(a). In particular, the Examiner has rejected claims 23-27 as being unpatentable over Talagala et al. (U.S. Patent Pub. No. 2002/0161972, hereinafter "Talagala"). Applicant respectfully traverses the rejection.

Talagala fails to teach, disclose, motivate or suggest method for recovering data as recited in Applicant's independent claim 23. In particular, Applicant's independent claim 23 recites:

"A method for recovering data from a failed disk drive in a disk array comprising a plurality of disk drives, said plurality of disk drives having a plurality of extents apportioned equally into a plurality of parity groups, said method comprising:

striping data from at least one file sequentially across the plurality of extents of the plurality of parity groups, said striping data comprising dividing each extent in each parity group into a plurality of data segments, storing data information in the plurality of data segments, forming a plurality of parity segments from the data segments in the parity group, and appending one parity segment to the end of each extent, wherein parity segments in the parity group respectively store parity information corresponding to the parity group in which the parity segments reside;

identifying a failed parity group containing an extent associated with the failed disk drive;

279205-1

Serial No. 09/940,728
Page 22 of 25

outputting, in realtime to users in a normal disk access mode, said data from parity groups without the failed disk drive;
reconstructing, in a parity correction mode of operation, said data from the at least one failed parity group; and
outputting, in realtime to the users in the failed parity group, said reconstructed data." (emphasis added).

The test under 35 U.S.C. §103 is not whether an improvement or a use set forth in a patent would have been obvious or non-obvious; rather the test is whether the claimed invention, considered as a whole, would have been obvious. Jones v. Hardy, 110 USPQ 1021, 1024 (Fed. Cir. 1984) (emphasis added). Moreover, the invention as a whole is not restricted to the specific subject matter claimed, but also embraces its properties and the problem it solves. In re Wright, 6 USPQ 2d 1959, 1961 (Fed. Cir. 1988) (emphasis added). Talagala fails to teach or suggest the Applicant's invention as a whole.

In particular, Talagala fails to teach or suggest the feature of "striping data from at least one file sequentially across the plurality of extents of the plurality of parity groups, said striping data comprising dividing each extent in each parity group into a plurality of data segments, storing data information in the plurality of data segments, forming a plurality of parity segments from the data segments in the parity group, and appending one parity segment to the end of each extent, wherein parity segments in the parity group respectively store parity information corresponding to the parity group in which the parity segments reside."

Talagala discloses "data and parity are striped across the storage devices 1 through 5. The data stripe corresponding to data and parity blocks for the A data blocks are the same as that shown in FIG. 3. When processor 100 of FIG. 1 writes new data to array of storage devices 410 of FIG. 4, the data is again striped across the storage devices. In this example, data stripe "B" represents new data written to array of storage devices 410. The data is broken into four blocks (stripe units), B(0) through B(3) and a parity block P(B) is calculated. The data blocks B(0) through B(3) and P(B) are stored across the storage devices such that the data and parity blocks are not stored on the same storage device." (see Talagala, page 4, paragraph 0043, and FIGS. 3 and 4). In other words, each block of data is respectively stored as a stripe unit across the array of

279205-1

Serial No. 09/940,728
Page 23 of 25

disk drives (see Talagala, page 3, paragraph 0033 (i.e., "the data has been divided into four data blocks A(0) through A(3) and stored on storage devices 1 through 4, respectively.").

Talagala fails to teach or suggest dividing each extent in each parity group into a plurality of data segments. Specifically, the blocks of data illustratively represented by blocks B(0) through B(3) each represent an extent of data that is stored as a stripe unit on each disk drive, as opposed to an extent being divided into a plurality of data segments, which are collectively stored as a stripe unit. Thus, each stripe unit shown and discussed with respect to FIGS. 3 and 4 of Talagala represents an extent striped across a disk drive. In other words, each stripe unit is an extent, as opposed to being a plurality of segments forming an extent as claimed by the Applicant's invention.

Referring to the Applicant's specification:

FIG. 4 depicts a RAID 3+5 data striping format across a plurality of extents in the disks 320 in a single parity group 404. The parity group 404 is illustratively depicted as having four data extents 406 (extents 0-3) containing data (e.g., video information). Each data extent 406 is divided into segments 402₁ to 402₃ (collectively segments 402). The number of segments 402 in each extent containing data is equal to P-1, where P equals the number of disks 320 in a parity group 404. For example, the true (non-parity) data of extent 0 110₀ is divided into three segments (a, b and c). The true data of extent 1 110₁ is divided into segments d, e, and f. The true data of extent 2 110₂ is divided into segments g, h, and i, and the true data of extent 3 110₃ is divided into segments j, k, and l.

The parity column 408 is a representation of the parity information for the segments as the data is stored across the parity group 404. For instance, the parity data "X" for the first row of extent segments 402₁ includes the data in segments a, d, g, and j of extents 0-3, respectively. Similarly, the parity data "Y" for the second row of extent segments 402₂ includes the data in segments b, e, h, and k of extents 0-3, respectively. Furthermore, the parity data "Z" for the third row of extent segments 402₃ includes the data in segments c, f, i, and l of extents 0-3, respectively. It should be noted that the parity information in column 408 is not a physical disk drive 320 that stores the parity information. Rather, column 408 is simply a representation of the parity information for the four data extents 406 of the parity group 404.

At the bottom of extent 0 110₀ is appended the parity segment X 410₀, which contains the parity data of the first segments 402₁ of each extent ($X = a \wedge d \wedge g \wedge j$) in the parity group 404, where " \wedge " indicates a Boolean logic "exclusive OR" (XOR) operative. Similarly, at the bottom of extent 1 110₁ is appended the parity segment Y 410₁, which contains the parity data of the second segments 402₂ of each extent ($Y = b \wedge e \wedge h \wedge k$) in the parity group 404. Likewise, at the bottom of extent 2 110₂ is appended the parity segment Z 410₂, which contains

279205-1

Serial No. 09/940,728
Page 24 of 25

the parity data of the third segments 402₂ of each extent ($Z = c^{f^4}$) in the parity group 404. Furthermore, at the bottom of the last extent 110₃ (i.e., extent 3) is appended the parity segment "Q" 410₃, which contains the parity information of the 3 parity segments X, Y, and Z, where $Q = X^Y^Z$. Thus, in the inventive RAID 3+5 format, the "pure parity" disks required for RAID 3 (see FIG. 1) are eliminated from the disk drive array 319. (see Applicant's specification, pages 9 and 10, paragraphs 0033, 0034, and 0036 and FIG. 4).

Thus, each disk drive is divided into a plurality of extents, where each extent is capable of storing a single block of data. The Applicant's invention further divides each extent in each disk drive associated with a parity group into a plurality of segments containing data, and appends a parity segment to the end of each extent, where the parity information corresponds to the data stored in the data segments. The Applicant's invention is completely different from the Talagala reference, since there is no teaching or suggestion of dividing each extent in each parity group into a plurality of data segments. Therefore, Talagala fails to teach or suggest the Applicant's invention as a whole.

Furthermore, claims 25-27 depend directly or indirectly from independent claim 23 and recite additional features thereof. As such and for at least the same reasons as discussed above, the Applicant submits that these dependent claims are not obvious and fully satisfy the requirements of 35 U.S.C. §103. Therefore, Applicant respectfully submits that the Examiner's rejection of claims 25-27 should be withdrawn.

Serial No. 09/940,728
Page 25 of 25

CONCLUSION

Thus, Applicant submits that all of the claims presently in the application, are enabling, not anticipated, non-obvious and patentable under the respective provisions of 35 U.S.C. §112, §102 and §103. Accordingly, both reconsideration of this application and its swift passage to issue are earnestly solicited.

If, however, the Examiner believes that there are any unresolved issues requiring adverse final action in any of the claims now pending in the application, it is requested that the Examiner telephone Eamon J. Wall, Esq. or Steven M. Hertzberg at (732) 530-9404 so that appropriate arrangements can be made for resolving such issues as expeditiously as possible.

Respectfully submitted,

Dated: _____

8/25/04

EJ Wall

Eamon J. Wall
Registration No. 39,414
Attorney for Applicant

MOSER, PATTERSON & SHERIDAN, LLP
595 Shrewsbury Avenue, Suite 100
Shrewsbury, New Jersey 07702
Telephone: 732-530-9404
Facsimile: 732-530-9808